

Performing Survival Analysis on HPC System Memory Error Data

Author: Stephen Penton

Mentors: Dr. Nathan DeBardeleben & Terry Grové

Statistical analysis of time-to-event variables presents a unique challenge due to their nature. As traditional regression techniques are not sufficient to fully capture time-to-event information, survival analysis is a branch of statistics used to answer questions about a population's lifetime including the rate at which individuals experience specific events. These techniques require specific data and there are various models that can be utilized to perform different analyses. We present a tool that ingests data from a user and is able to produce an initial set of results for survival analysis at varying levels of complexity. We demonstrate its capabilities by testing it on system memory error data of two LANL machines: Cielo and Trinitite. We show the tools use within an HPC specific domain to provide insight into the impact of specific features on when memory errors occur.